



HPSS at SDSC

Joseph Lopez

lopez@sdsc.edu



NATIONAL PARTNERSHIP FOR ADVANCED COMPUTATIONAL INFRASTRUCTURE

SAN DIEGO SUPERCOMPUTER CENTER



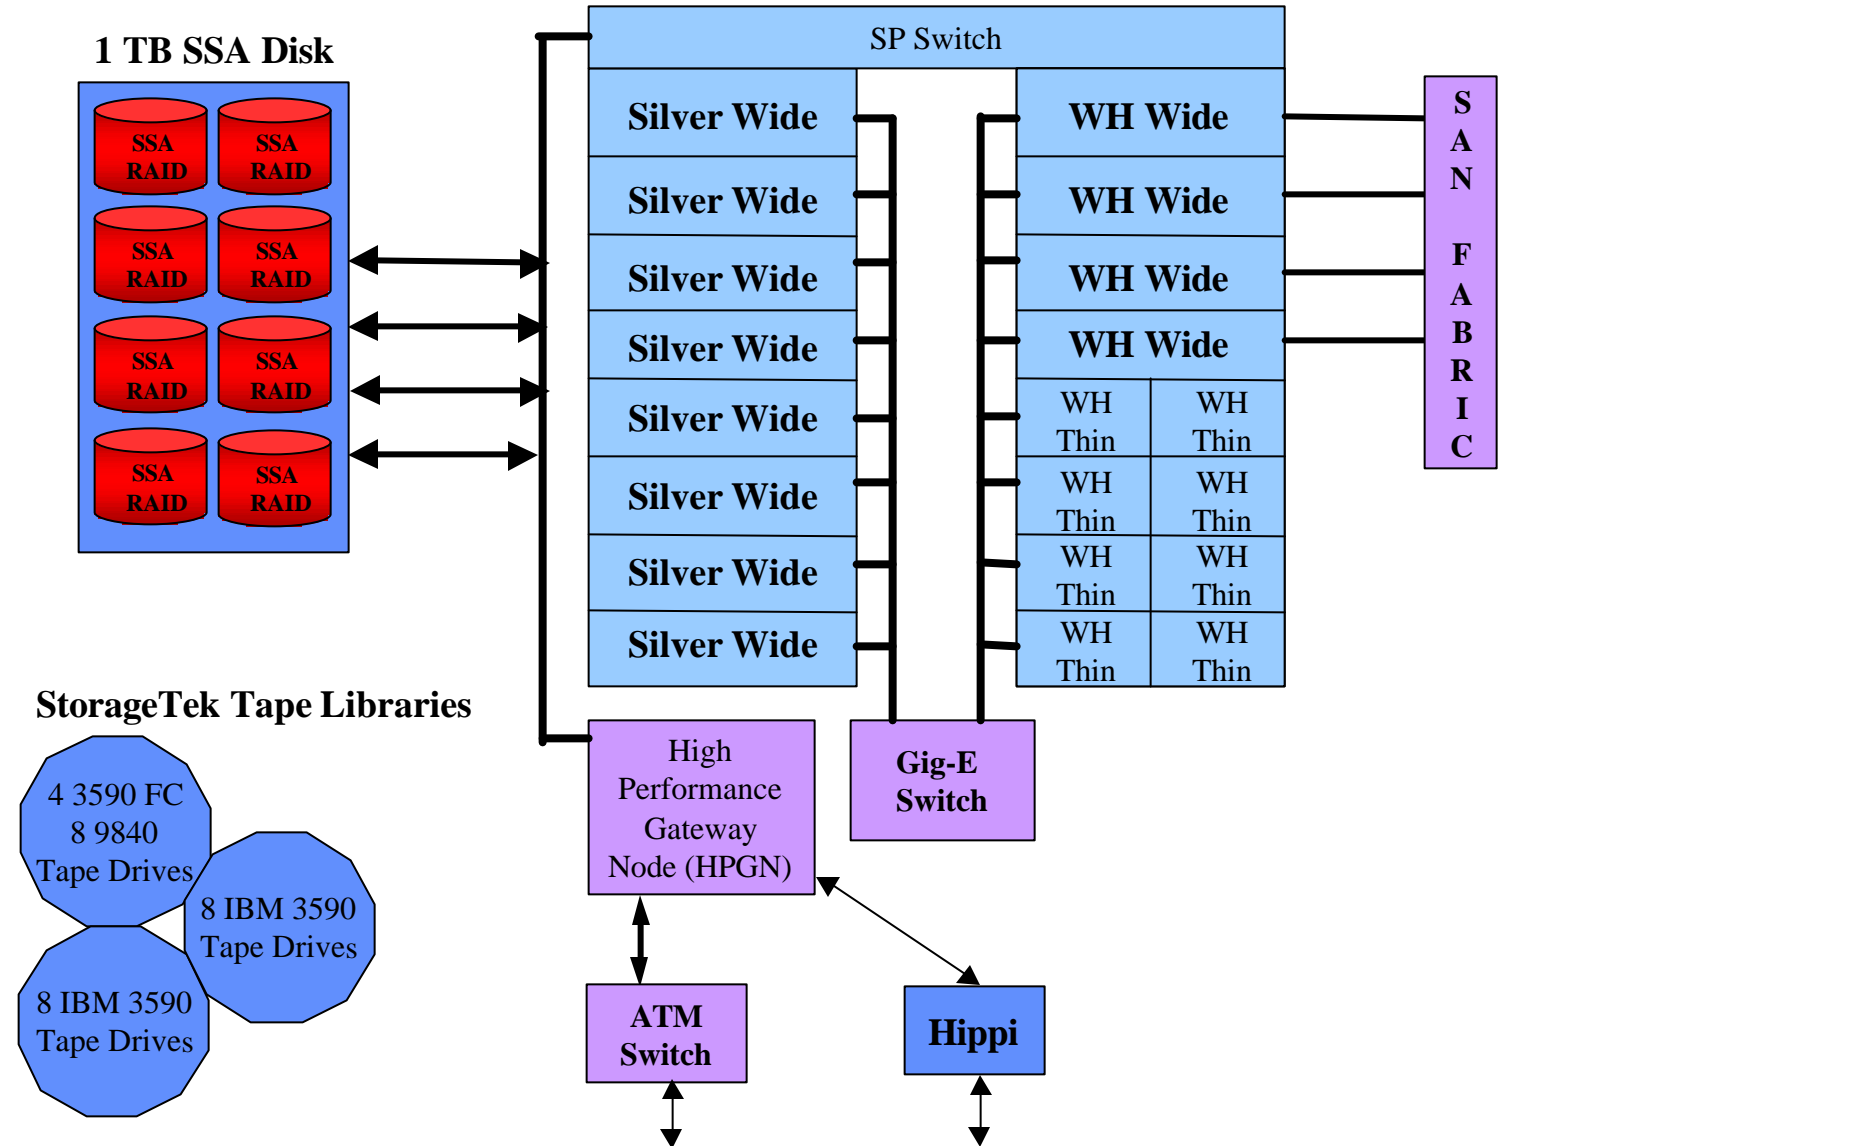
HPSS Hardware Configuration

- 20 SP nodes
 - 4 Silver wide
 - 4 Winterhawk wide
 - 8 Winterhawk thin
 - Gigabit ethernet
 - Hippi
 - ATM via HPGN

Configuration

- 2 TB disk cache
 - 1 TB SSA raid
 - 1 TB Sun T3 fiberchannel raid in production
 - Up to 3 TB possible
 - 3 STK Powderhorns
 - 28 tape drives
 - 20 IBM 3590E
 - 8 STK 9840

Mass Storage (HPSS) Configuration



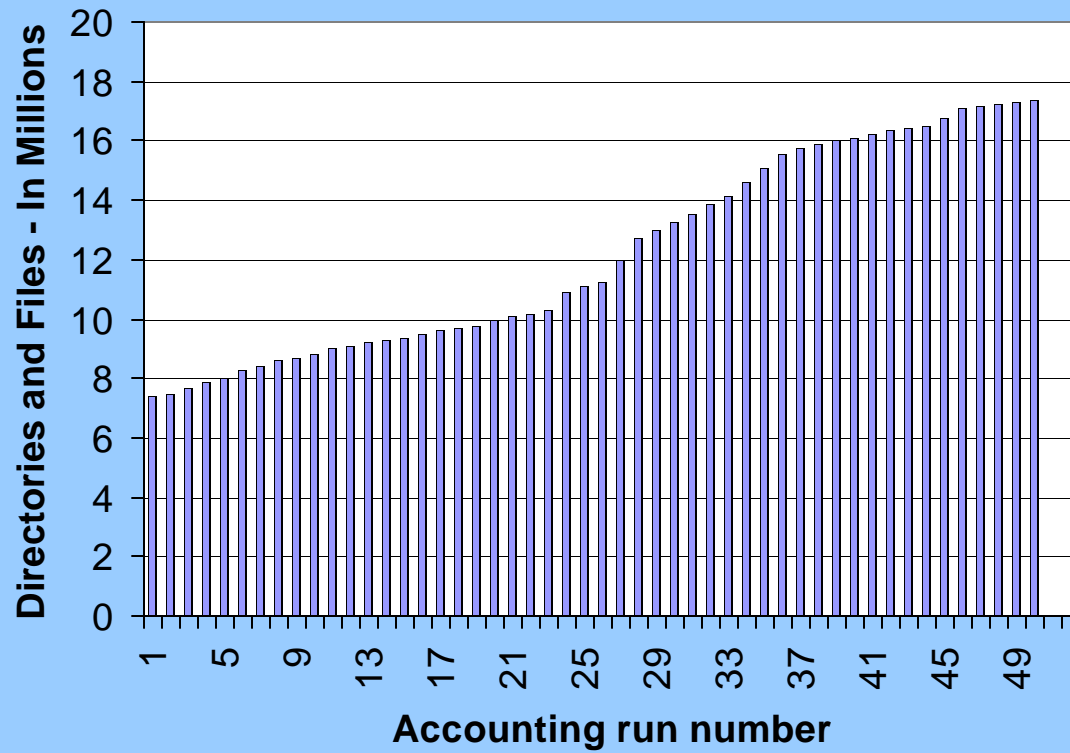
Software Level

- AIX 4.3
- DCE 2.2
- Encina 4.2
- HPSS 4.1.1
- HPSS user interfaces
 - Hsi
 - Pftp
 - SRB

Data Storage

- 280 TB Stored
- 14.8 million files - 3 million directories
 - NPACI and PACI research applications
 - System backups
 - Industrial users
 - User Workstations
 - 11,373 users

HPSS Growth



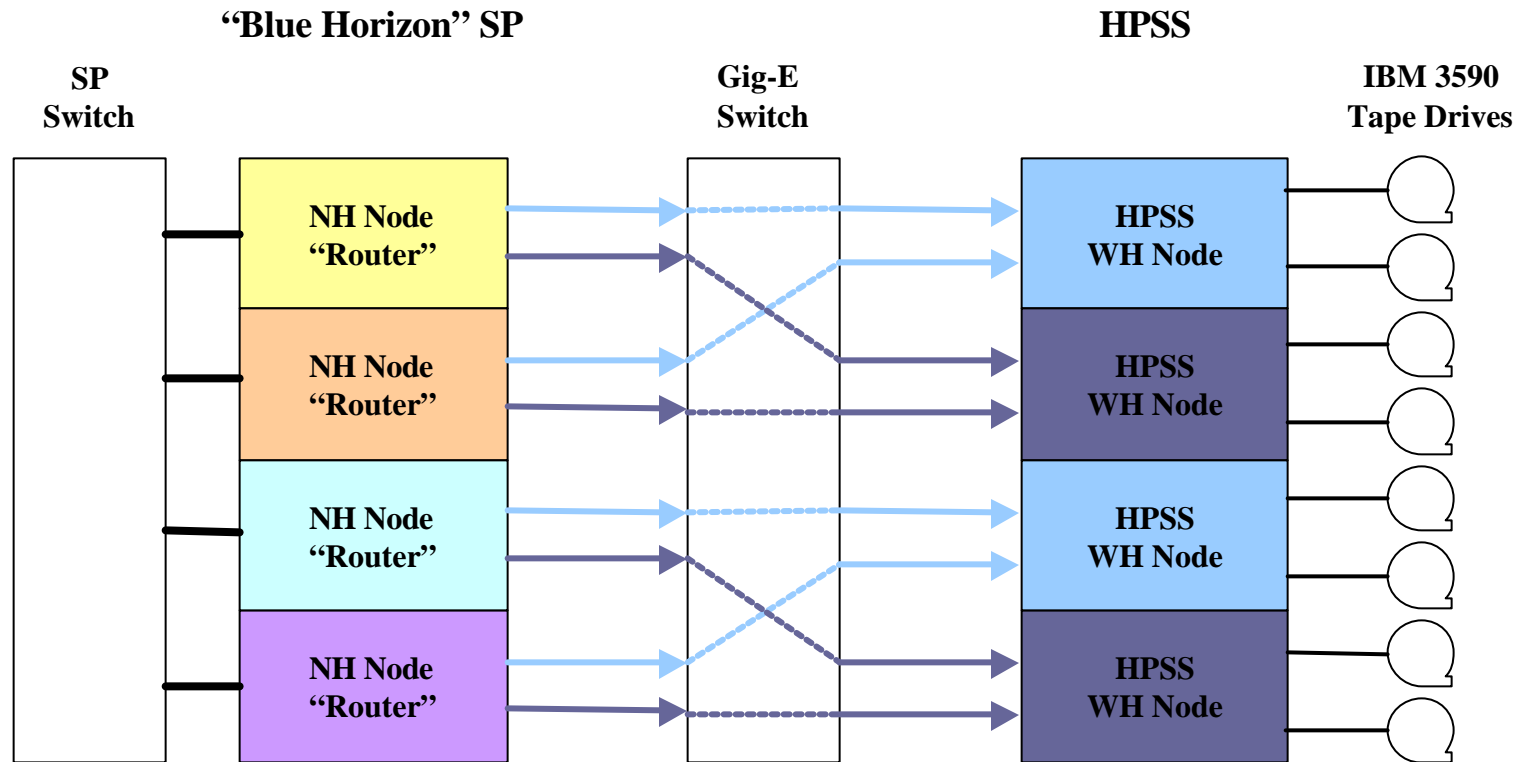
To do list

- Upgrade to 4.2 as soon as possible
- Encourage NPACI and PACI sites to use Hsi
- Repack all 3590 to 3590E tape
- Add more tape storage and drives
- Improve performance
 - Reconfigure Storage classes
 - Retire one copy small file cos
 - Consolidate small volumes

HPSS to T3 tests

- 4 Brocade 2800 16 port switches
- 14 Sun T3 bricks (36G and 72G)
- 4 IBM 3590E tape drives
- 10 Sun hosts
- 12 IBM SP hosts


8 way striped data transfers to HPSS



Striped Data Transfers to HPSS

- Machines are logically sub-netted at the router
- “Blue Horizon” SP organized in 4 network ‘quadrants’
- HPSS servers divided across 2 networks
- “Router” nodes do network I/O to HPSS on behalf of remaining nodes in a quadrant.

Transfer Rates: Blue Horizon to HPSS Disk



Issues and Concerns

- Tape read problems
 - Bad tapes propagate through all drives
 - Relabeling and imports very slow

Wish list

- Incorporate Hsi and Htar into HPSS
- Add comment field to devices and drives when disabling for administrative purposes
- Have Houston run a production system onsite for reproducing reported problems
- More useful content in messages where appropriate
 - for example - STK responses